Topological Optimisation of Rod-Stirring Devices

Matthew D. Finn*

School of Mathematical Sciences, University of Adelaide, Adelaide SA 5005, Australia

Jean-Luc Thiffeault†

Department of Mathematics, University of Wisconsin, Madison, WI 53706, USA

(Dated: April 6, 2010)

There are many industrial situations where rods are used to stir a fluid, or where rods repeatedly stretch a material such as bread dough or taffy. The goal in these applications is to stretch either material lines (in a fluid) or the material itself (for dough or taffy) as rapidly as possible. The growth rate of material lines is conveniently given by the topological entropy of the rod motion. We discuss the problem of optimising such rod devices from a topological viewpoint. We express rod motions in terms of generators of the braid group, and assign a cost based on the minimum number of generators needed to write the braid. We show that for one cost function—the topological entropy per generator—the optimal growth rate is the logarithm of the golden ratio. For a more realistic cost function, involving the topological entropy per operation where rods are allowed to move together, the optimal growth rate is the logarithm of the silver ratio, $1 + \sqrt{2}$. We show how to construct devices that realise this optimal growth, which we call *silver mixers*.

arXiv:1004.0639v1 [nlin.CD] 5 Apr 2010

I. INTRODUCTION

A rod-stirring device, where a number of rods are moved around in a fluid, is the most natural and intuitive method of stirring. The number of rods, their shape, and the nature of their motion constitute a stirring protocol. For example, moving a spoon in a figure-eight pattern in a mixture is a simple way of blending the ingredients of a cake. But beyond applications in the kitchen, such stirring methods are widely used in industry. For instance, when glass is manufactured, it is in a molten state. It is inhomogeneous both in temperature and composition. These inhomogeneities are undesirable: the human eye can detect very small variations in density, and modern glass must be of the highest quality. One way to remove these inhomogeneities is to stir the molten glass before it cools.

There, problems begin: you cannot stir molten glass the way you would a cup of coffee—it is simply too viscous. Hence, the stirring rods are limited in their speed and path, and it is important to know what paths are optimal in order to improve the performance of a mixing device. This is where mathematics can help, and is the topic of this paper: we introduce a topological approach to the optimisation of rod-stirring devices. The topological approach was pioneered for fluids by Boyland *et al.* [1] and developed by many others [2–17]. Of course, glass mixing is only a representative application, and the work presented

*Electronic address: matthew.finn@adelaide.edu.au

†Electronic address: jeanluc@math.wisc.edu

here applies to many other important forms of mixing, such as industrial dough production which uses a type of device called a pin mixer [18] (see Fig. 1). In many cases, as in taffy-making (see Fig. 2), there is not necessarily an underlying 'fluid' filling the space, but the optimisation techniques we discuss apply directly and so we use the same language to describe these devices.

In all physical problems, modelling begins by making simplifying assumptions. For glass mixing and for many other situations, the fluid motion is predominantly two-dimensional. This is due in part to the thinness of the glass layer, but also to its thermal stratification. Hence, it is suitable and highly desirable to model the system as an idealised two-dimensional fluid. Mathematically, this means that the domain of fluid motion is a two-dimensional space—a surface. More precisely, it is a punctured surface with an outer boundary, since the stirring creates obstacles in the space around which the fluid must flow. Hence, if we are only interested in a topological characterisation of the mixing region, we can consider our simplified mixing device to be the punctured disc, where the movable punctures are the stirring rods. By topological characterisation, we mean that we do not care about the specific size of features, only about their global impact on the space.

The fluid motion induced by any periodic motion of the rods (or punctures) defines a mapping, that is, a rule for describing how small parcels of fluid are dragged by the rods at each stirring cycle. This mapping is normally obtained by solving the relevant fluid equations: the Stokes equations for highly-viscous fluids such as molten glass, or more complicated equations when stirring a polymeric fluid. But the beauty of a topological approach is that we can deduce much about the range of possible outcomes of rod motion without ever specifying governing equations of motion. From this point of view, it suffices that the fluid is an idealised continuum, an assumption that holds for essentially all fluids (and even granular materials in some limits).

Hence, the mathematical setting for our problem is the space of all possible mappings of the punctured disc that arise from rod motions. This is a gigantic space, and it contains as special cases all possible fluid motions. The key to making sense of this space is that there is a theorem—the Thurston–Nielsen classification theorem [19, 20]—that tells us that the members of that space (the mappings induced by rod motions) can be grouped into three (and only three!) categories.

The first category is called finite-order. If we move the rods around following this type of motion, and return them to their original position, then in a well-defined sense material lines in the fluid have not become 'entwined' on the rods. (These mappings are isotopic to the identity.) Since our ultimate goal is to mix the fluid by stirring, this is terrible. It means that we are wasting our precious time and energy with this particular rod motion. We shall therefore say nothing more about this category.

The second category of mappings is called pseudo-Anosov, and it is much more interesting. A pseudo-Anosov mapping leads to a complex intertwining of material lines in the fluid around the rods. In fact, a typical material line is forced to grow exponentially under repeated stirring. The pseudo-Anosov mappings are closely related to chaotic dynamical systems. They are the best type to stir with, and for appropriate rod motions they can be very effective. This was demonstrated by Boyland *et al.* [1], who constructed a rod-stirring device with a pseudo-Anosov rod motion. (A mechanically more straightforward implementation was given by Binder and Cox [2].) Their design is, however, far from optimal, and we will improve upon it in this paper. Note that to have a pseudo-Anosov mapping we need at least three rods: it is not possible to have topological complexity with only one or

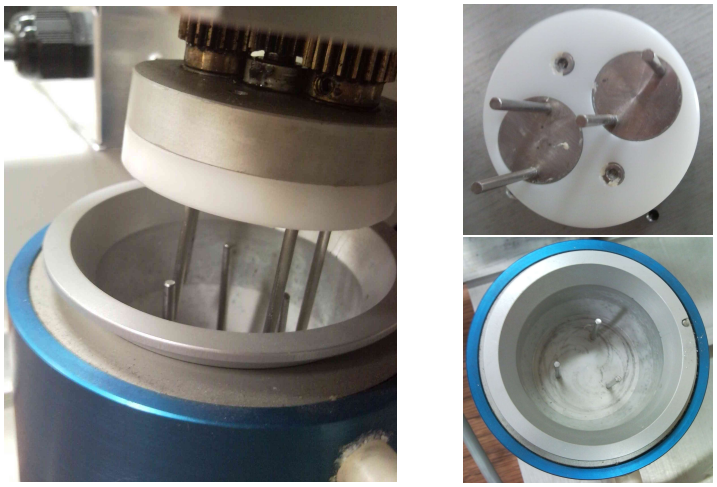

(a) (b)

FIG. 1: (a) The mixograph, a model planetary pin mixer for bread dough. (b) Top section with four moving rods (above), and bottom section with three fixed rods (below). See Connelly and Valenti-Jordan [18]. (Courtesy of the Department of Food Science, University of Wisconsin.)

two rods [1], unless artificial 'punctures' are made in the space [4]. This helps explain why the taffy puller in Fig. 2 has three rods.

The third category of mappings is called reducible. Essentially, mappings in that category break up the fluid domain into separate regions, where each region belongs to the first or second categories above. Mappings in this category are likely to have barriers to mixing, which is undesirable. We shall not discuss maps of this type any further.

Our aim in this paper is to select rod motions that belong to the second category—pseudo-Anosov—and are particularly good at mixing. To accomplish this, we first need a convenient manner of describing rod motions topologically. We show how to do this in terms of braids in Section II. Next, we need to specify our measure of 'good mixing,' the topological entropy: this is the minimum complexity imparted on the fluid trajectories based on the rod motions (see Section III). We also need to ascribe a cost to the rod motions, and we present two possible choices in Section IV. One of these measures of cost we argue the more physically relevant. In Sections V and VI we describe a rod-stirring device that incorporates our optimality criterion, and show results from simulations and experiments. We offer some concluding remarks in Section VII.

II. ROD MOTIONS AS BRAIDS

In the introduction, we saw that the mapping describing the fluid motion belongs to one of three categories: finite-order, pseudo-Anosov, or reducible. Only the second is important for our purposes. But even within the category of pseudo-Anosov mappings, there are many *equivalence classes* of mappings. This means that there are mappings that are fundamentally different, in that they cannot be obtained from one another by a continuous transformation. This concept is called *isotopy* of mappings, and the set of all mappings isotopic to a given mapping is called its *isotopy class*. For a more detailed discussion see [1, 3–9].

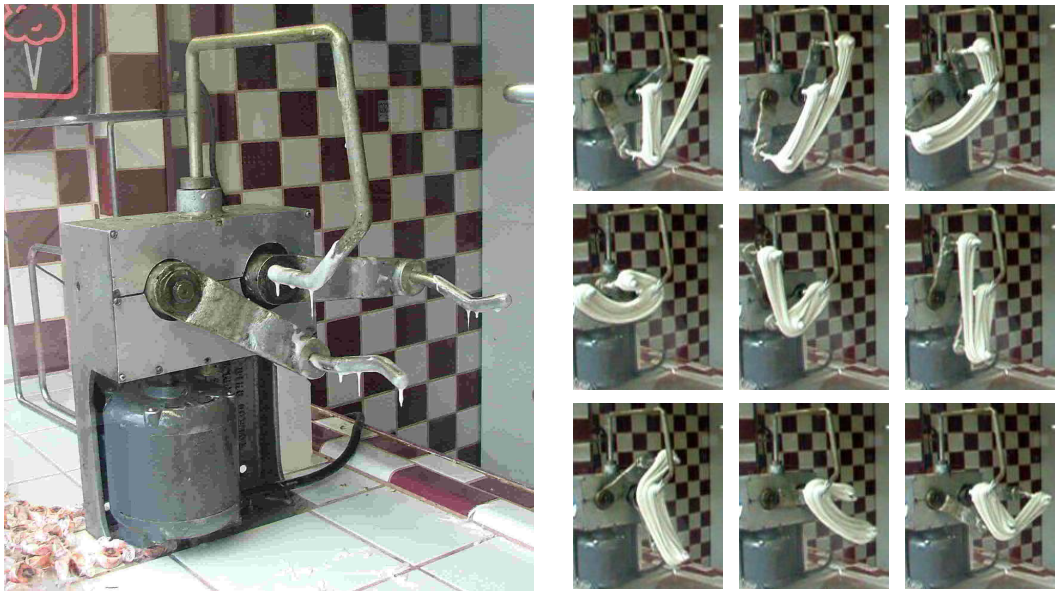

(a) (b)

FIG. 2: (a) A taffy-pulling device. (b) Snapshots over a full period of operation, with taffy.

Now we introduce a convenient way to describe rod motions: braids. Picture a given two-dimensional rod motion as a three-dimensional space-time plot, as in Fig. 3(a), where the vertical axis represents time. The resulting tangle of strands is called a *physical braid*. Two physical braids are equivalent if they can be deformed into each other with no strands crossing other strands or boundaries.

An example of a device whose rod motion can be described by a physical braid is shown in Fig. 2(a), which is a photograph of a *taffy puller*. (The purpose of the taffy puller is to stretch and fold candy to oxygenate it, making it lighter and chewier.) The central rod is fixed, and two rods orbit the central one following a motion shown in Fig. 2(b) and as a space-time plot in Fig. 3(a). Fixed rods are often called *baffles*, though the distinction between fixed and moving rods is immaterial to a topological description. We will analyse the taffy puller in more detail in the next section.

The physical braid describing the rod motion of the mixograph in Fig. 1 is a little more complicated. The motion may be viewed in a frame of reference where either the three lower baffles are fixed (Fig. 4(a)) or are rotating (Fig. 4(b)). The braids appear different in both cases, but yield the same topological information about the complexity of the flow.

We can pass from physical braids to an algebraic description of braids by introducing *generators*. Assume without loss of generality that all n stirrers (*i.e.*, stirring rods) can be ordered from left to right along some projection line. The generator σ_i, $i = 1, \ldots, n-1$, denotes the clockwise interchange of the ith rod with the $(i+1)$th rod along a circular path, where i is the position of a rod counting from left to right. The anticlockwise interchange is denoted σ_i^{-1} (see Fig. 3(b) for a depiction of generators.) We can write consecutive interchanges as a sequence, for example $\sigma_1\sigma_2^{-1}$, where the generators are read from left to right (this convention differs from Boyland *et al.* [1]). This composition law allows us to define the *Artin braid group on n strands*, B_n, with the identity given by untangled strands.

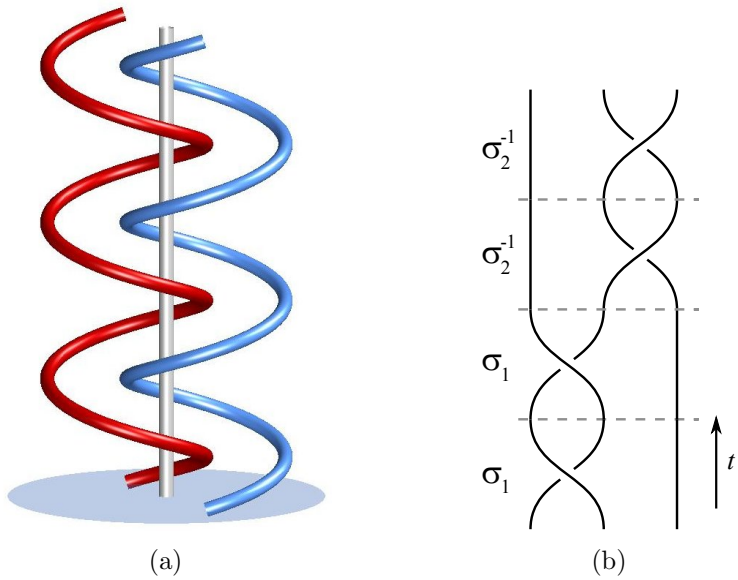

FIG. 3: (a) The trajectories of the three rods of the taffy puller plotted in a space-time diagram, with time flowing from bottom to top. This defines a braid on $n = 3$ strands, which in terms of braid group generators is written $\sigma_1^2\sigma_2^{-2}$, read from left to right (three periods are shown). (b) The diagram for the braid $\sigma_1^2\sigma_2^{-2}$, which has dilatation $(1+\sqrt{2})^2$.

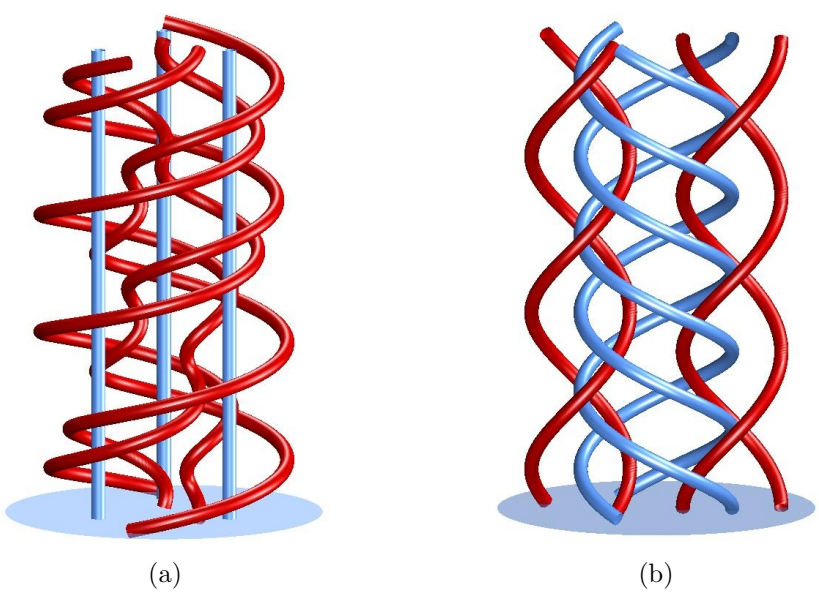

FIG. 4: (a) The trajectories of the seven rods of the mixograph in Fig. 1 viewed in a frame of reference where the three lower baffles are fixed. (b) The same motion, but viewed in a rotating frame where the four upper pins rotate in circles and the lower mixing vat and baffles counter-rotate in the opposite direction. (Three periods are shown.) In (b) the rod paths repeat the $n = 7$ strand braid $\sigma_3\sigma_2\sigma_3\sigma_5\sigma_6^{-1}\sigma_2\sigma_3\sigma_4\sigma_3\sigma_1^{-1}\sigma_2^{-1}\sigma_5$. The dilatation associated with this braid is approximately 4.1858, and its topological entropy is 1.4317.

A sequence of generators is thus an element of B_n and is called a *word* or *braid word*. (Repeated generators are written as powers, such as $\sigma_1\sigma_1 = \sigma_1^2$ and $\sigma_2^{-1}\sigma_2^{-1} = \sigma_2^{-2}$.) Braids are often written in a standard form as a *braid diagram*, as in Fig. 3(b).

In order that braid group elements correspond to physical braids, they must satisfy the additional relations [21]

$$\sigma_i\sigma_j\sigma_i = \sigma_j\sigma_i\sigma_j \qquad \text{if} \qquad |i-j| = 1, \tag{1a}$$

$$\sigma_i\sigma_j = \sigma_j\sigma_i \qquad \text{if} \qquad |i-j| \geq 2. \tag{1b}$$

These relations are sufficient to characterise B_n, a nontrivial fact first proved by Artin [22].

The crucial role of braids as an organising tool is that they correspond to isotopy classes. (There is a subtlety involving possible rotations of the outer boundary, so that one usually speaks of *braid types* specifying isotopy classes rather than braids [1, 8].) We assign a braid to stirrer motions by the diagrammatic approach employed in Fig. 3, and to every braid we can assign stirrer motions by reversing the process. Exploring possible stirring protocols is now reduced to the algebraic task of examining braids, which is more suitable to optimisation.

III. COMPUTING THE TOPOLOGICAL ENTROPY OF A BRAID

As we saw in Section II, a periodic rod-stirring protocol can be associated with a braid on n strands, where n is the number of rods. Corresponding to this braid is a very useful measure of mixing and complexity, the *topological entropy* of the braid. The topological entropy of a braid is the minimum stretching rate that the rod motion imparts on material lines in the flow [1, 5, 23–25]. This minimum stretching rate is due to material lines being dragged along by the rods, since they are not permitted to go through the rods (think of the children's game 'cat's cradle' [26], where fingers play the role of rods). If a braid induces a mapping that is in a pseudo-Anosov mapping class, then the topological entropy is positive. This leads to material lines growing exponentially, which is good for mixing as it implies the interface between two solutes becomes more and more convoluted with time. The topological entropy is always non-negative, and it forces exponential stretching on material lines only if it is strictly positive. An equivalent measure is the *dilatation* of a braid, which is simply the exponential of the entropy, and conversely the entropy is the logarithm of the dilatation.

The topological entropy is only a lower bound on the rate of stretching of material lines: the lines can develop convoluted 'secondary folds' which are not forced by rods. This occurs, for instance, if the material lines are not pulled tight around the rods, but rather have loose sections that fold upon themselves. In fluid dynamics there is usually a gap between the lower bound and the measured rate of stretching of material lines in the fluid. In applications such as bread dough mixing and taffy pulling the lower bound is typically sharp, since the elastic nature of the material tends to make it wind tightly around the rods.

The problem is then to compute the topological entropy and dilatation of a given braid. For three strands (*i.e.*, three rods), the Burau representation offers an efficient method of computing the entropy [27–30]. The Burau representation associates with each generator a two-by-two matrix,

$$[\sigma_1] = \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix}, \qquad [\sigma_2] = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \tag{2}$$

where the square brackets indicate the matrix representation of a generator. The representation for σ_1^{-1} and σ_2^{-1} is obtained by matrix inversion. The Burau representation of a braid

word is then given simply by multiplying the corresponding matrices together. The spectral radius (the magnitude of the largest eigenvalue) of the Burau representation of a braid word then gives the dilatation of the braid [27–30]. For example, for the braid word $\sigma_1^{-1}\,\sigma_2$, the Burau representation is

$$[\sigma_1^{-1}\,\sigma_2] = [\sigma_1^{-1}]\cdot[\sigma_2] = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}. \tag{3}$$

This matrix has spectral radius $(3+\sqrt{5})/2$, which is the dilatation, and hence the entropy is $\log[(3+\sqrt{5})/2]$. (As we will discuss later, this dilatation is the square of the golden ratio $(1+\sqrt{5})/2$.) Roughly speaking, the Burau representation counts the number of windings of loops around the rods, so its spectral radius measures the growth of these loops under repeated application of the map.

As another example, we return to the taffy puller (Section II) which executes the braid $\sigma_1^2\sigma_2^{-2}$. The Burau representation is

$$[\sigma_1^2\sigma_2^{-2}] = [\sigma_1]^2\cdot[\sigma_2^{-1}]^2 = \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix}^2 \cdot \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}^2 = \begin{pmatrix} 1 & -2 \\ -2 & 5 \end{pmatrix},$$

which has the spectral radius $3+2\sqrt{2} = (1+\sqrt{2})^2$, as presented above.

Computing the topological entropy using the Burau representation is the simplest approach for three rods ($n = 3$). For more than three rods, however, this approach only provides a lower bound. Moreover, this lower bound is often nowhere near sharp. Hence, a reliable computation of the entropy for $n > 3$ requires a different approach. The most powerful algorithm is due to Bestvina and Handel [31]: for a given braid, the algorithm gives its isotopy class (*i.e.*, finite-order, pseudo-Anosov, or reducible) as well as its entropy. The algorithm is fairly complex, and a discussion of it is outside the scope of this paper, but there is a C++ implementation written by Toby Hall [32] that we use here. When speed is important (as when searching for optimal braids), we use the iterative algorithm of Moussafir [16], also described in Thiffeault [7].

IV. OPTIMISING OVER GENERATORS

We have introduced in Section II the idea of representing two-dimensional rod motions by a braid on n strands, where n is the number of rods. Then in Section III we discussed the dilatation and topological entropy of a braid, which are lower bounds on the amount of stretching experienced by material lines wrapped around the rods in each stirring period. We now turn to the question of optimisation: clearly some stirring protocols are better than others, and we would like to know which ones are best. We will take a topological viewpoint, assuming that topological entropy is the relevant quantity to optimise. The other part is choosing a cost function for our optimisation, and for this there are several choices. We shall look at two in particular: topological entropy per generator (TEPG), and topological entropy per operation (TEPO), and find the latter to be better suited to practical applications.

A. Topological Entropy per Generator (TEPG)

Entropy can grow without bound as the length of a braid increases, so a proper definition of an optimal entropy requires a 'cost' associated with the braid. An obvious measure of

efficiency is to divide the entropy by the smallest number of generators required to write the braid word. For example, the braid $\sigma_1^{-1}\sigma_2$ has entropy $\log[(3+\sqrt{5})/2]$ and consists of two generators. Its topological entropy per generator (TEPG) is thus $\frac{1}{2}\log[(3+\sqrt{5})/2] = \log[(1+\sqrt{5})/2]$. The question of the maximal TEPG for a given braid group B_n is well-posed.

In B_3, D'Alessandro *et al.* [33] proved that the maximal TEPG is given by repeating the word $\sigma_1^{-1}\sigma_2$.[1] Thus, the maximal TEPG is $\log[(1+\sqrt{5})/2]$, the logarithm of the *golden ratio* $\phi_1 = (1+\sqrt{5})/2$. In general, the *metallic means*[2] are defined by $\phi_m := \frac{1}{2}(m+\sqrt{m^2+4})$ (see Appendix A). We now summarise the results for higher n (see Appendix B for the sketch of a proof, aimed at the specialised reader):

- For $n = 4$, the braid $\sigma_1^{-1}\sigma_2\sigma_3^{-1}\sigma_2$ has maximal TEPG of $\log\phi_1$, the same TEPG as for $n = 3$.
- For $n > 4$, all pseudo-Anosov braids in B_n have TEPG strictly less than $\log\phi_1$.

(We first presented these results in [5], which were also independently formulated by Moussafir [16].) In other words, the highest TEPG of $\log\phi_1$ is only achieved for $n = 3$ or $n = 4$. Intuitively, the optimal braids for $n = 3$ and $n = 4$ involve a tight combination of three or four strands to achieve their high entropy per generator. A braid with more strands requires extra generators to make it irreducible (and hence pseudo-Anosov), thereby reducing the TEPG.

For more than a few rods, the rod motions with optimal TEPG are not very relevant practically, but we note that in B_n we can get arbitrarily close to a TEPG of $\log\phi_1$ by, for example, repeating the braid $\sigma_1^{-1}\sigma_2$ a large number of times, followed by a few generators to make the braid irreducible (pseud-Anosov). The TEPG of such a braid will converge towards $\log\phi_1$ as more $\sigma_1^{-1}\sigma_2$ motions are added.

B. Topological Entropy per Operation (TEPO)

We saw in Section IV A that the maximum topological entropy per generator (TEPG) is equal to (for $n \leq 4$) or is uniformly bounded above (for $n > 4$) by $\log\phi_1$, where $\phi_1 = (1+\sqrt{5})/2$ is the golden ratio. But for applications, optimising the TEPG is not so useful. This is because each generator is counted separately, the assumption being that rods are moved sequentially according to the generators. In practice, however, it is better to move as many rods at the same time as possible. From the point of view of braid words, this means that we should count commuting adjacent generators as a single operation. For example, the braid $\sigma_1^{-1}\sigma_3^{-1}\sigma_2\sigma_4$ contains four generators, but the corresponding rod motions can be performed in two operations, since the two motions $\sigma_1^{-1}\sigma_3^{-1}$ can be performed together, followed by $\sigma_2\sigma_4$. In this viewpoint we count the 'cost' of the braid $\sigma_1^{-1}\sigma_3^{-1}\sigma_2\sigma_4$ as being two, not four. Thus we are led to examine the topological entropy per operation, or TEPO, which is a more physically relevant quantity to optimise than the TEPG.

---

[1] This is a similar result to Blondel *et al.* [34], who compute the 'joint spectral radius' of a collection of matrices using different methods.

[2] These are sometimes called *silver means*, but that name is easily confused with the silver ratio or second metallic mean, which we shall use extensively in this paper.

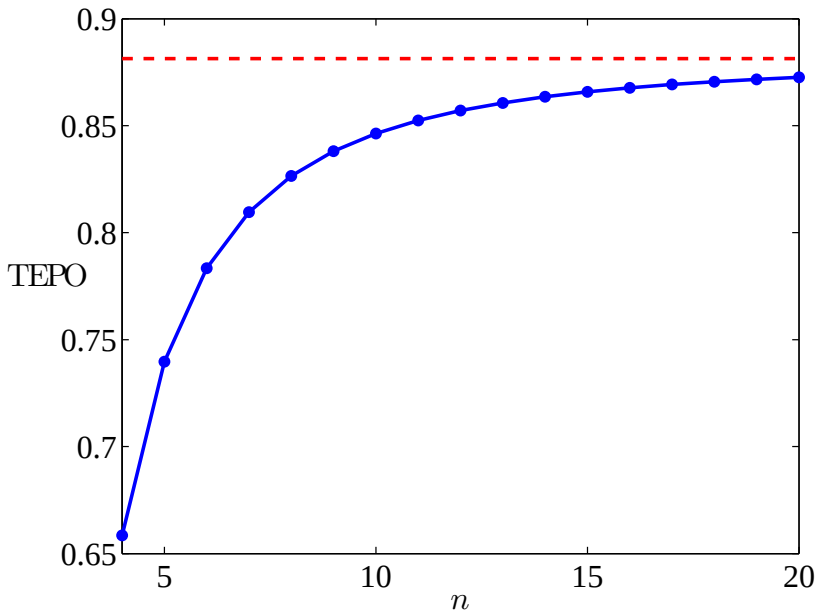

FIG. 5: Topological entropy per operation (TEPO) as a function of n, the number of stirring rods. The asymptote (dashed) is $\log \phi_2 \simeq 0.8814$, where $\phi_2 = 1 + \sqrt{2} \simeq 2.4142$ is the silver ratio.

Numerical investigation and the proof sketched in Appendix B suggest that the optimal TEPO is given by the braids

$$\sigma_1^{-1}\sigma_3^{-1}\cdots\sigma_{n-2}^{-1}\,\sigma_2\sigma_4\cdots\sigma_{n-1}, \qquad \text{for } n \text{ odd;} \tag{4a}$$

$$\sigma_1^{-1}\sigma_3^{-1}\cdots\sigma_{n-1}^{-1}\,\sigma_2\sigma_4\cdots\sigma_{n-2}, \qquad \text{for } n \text{ even.} \tag{4b}$$

The braids given in (4) consist of two operations: interchange anticlockwise the first and second rod, third and fourth, etc., then interchange clockwise the second and third rod, fourth and fifth, etc. These thus all have optimal TEPO because of the low number of operations required. In Fig. 5 we give the maximum TEPO as a function of n: notice that it is monotonically increasing, and it appears to asymptote to a fixed value for large n. In Section V and in Appendix B we show that the asymptotic value is $\log \phi_2$, where $\phi_2 = 1+\sqrt{2}$ is the *silver ratio*, or the second metallic mean (see Appendix A; recall from Section IV A that the first metallic mean ϕ_1 is the golden ratio).

C. Other Types of Optimisation

For practical applications, the type of optimisation we have discussed has limited applicability. The main issue is that the generators of the braid group, σ_i, do not necessarily correspond to natural motions of rods in a given physical system. For instance, the exchange of the first and last rod in a three-rod system is written $\sigma_1\sigma_2\sigma_1$ in terms of exchanges of neighbouring rods, so this simple operation requires three generators. But clearly it does not cost much more energy to do this than to exchange the first and second rods (σ_1). The generators do not capture the intrinsic geometry of the system. For $n = 4$, the exchange of the first and last rod requires five generators, making the apparent cost even greater.

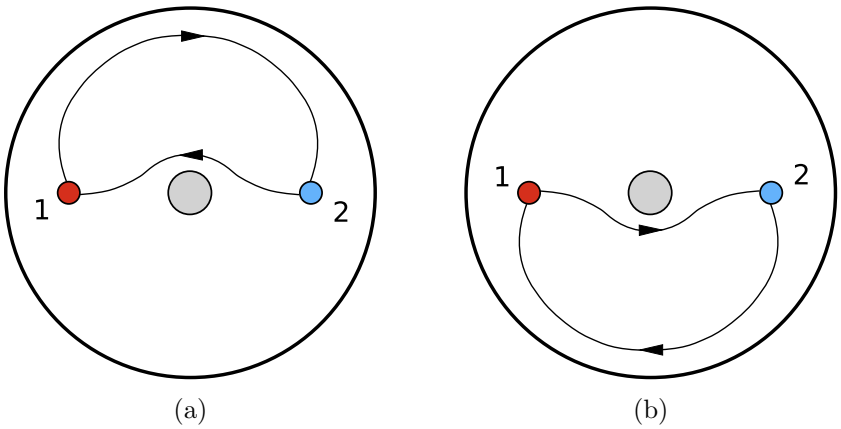

FIG. 6: Braid group generators for two rods in an annulus: (a) Σ_1; (b) Σ_2. The central baffle remains fixed.

Many different optimisation problems can thus be formulated to incorporate different engineering constraints. In particular, the resulting rod motions must be realisable using a straightforward design. We shall balance this engineering constraint with mathematical tractability in Section VI, where we discuss a class of mixing devices based on braids on an annulus.

V. BRAIDS ON AN ANNULUS

In Section IV B, we selected the topological entropy per operation (TEPO) as a suitable measure to optimise. We found that the optimal braid family given by (4) had TEPO that asymptotes to a fixed value for large n (Fig. 5). In the present section we will find the value of the asymptote by showing that the optimal braid for large n can be realised as a braid on an annulus.

Consider two moving rods in an annular geometry, as shown in Fig. 6. The central baffle is fixed and defines the annulus. We define the operation Σ_1 as the clockwise interchange of the two rods *above* the central baffle (Fig. 6(a)), and Σ_2 as the clockwise interchange of the two rods *below* the central baffle (Fig. 6(b)). These are the natural operations on an annulus with two rods. The generalisation to more rods is straightforward.

To compute the entropy of braid words in this annular domain with two rods and a central baffle, observe that we can also regard Fig. 6 as a three-rod system, in which case we can rewrite our two annular generators in terms of standard braid group generators as

$$\Sigma_1 = \sigma_1\,\sigma_2\,\sigma_1^{-1}\,, \qquad \Sigma_2 = \sigma_2\,\sigma_1\,\sigma_2^{-1}\,. \tag{5}$$

The standard braid group B_3 has the well-known Burau representation [27], given by Eq. (2). In Section III we explained that for $n = 3$ the spectral radius of the Burau representation of a braid word gives the dilatation of the braid. We now use the representation (2) to derive a matrix representation for (5),

$$[\Sigma_1] = [\sigma_1]\,[\sigma_2]\,[\sigma_1^{-1}] = \begin{pmatrix} 2 & 1 \\ -1 & 0 \end{pmatrix}, \qquad [\Sigma_2] = [\sigma_2]\,[\sigma_1]\,[\sigma_2^{-1}] = \begin{pmatrix} 0 & 1 \\ -1 & 2 \end{pmatrix}. \tag{6}$$

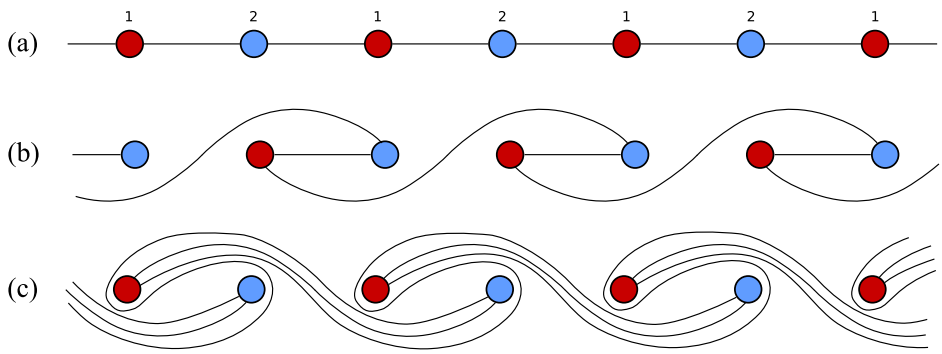

FIG. 7: (a) Two rods on a periodic lattice, joined by line segments. The segments are stretched by (b) the operation $H = \Sigma_2$, a clockwise interchange of the rods at positions 2 and 1; and (c) $V = \Sigma_1^{-1}$, an anticlockwise interchange of the rods at position 1 and 2.

The spectral radius of the representation of an annular braid word using (6) will thus give the exact dilatation of the braid. To make the connection between the annular braid group and the Artin braid group more explicit, we apply to (6) the coordinate transformation $[\Sigma_i'] = R\,[\Sigma_i]\,R^{-1}$, with

$$R = \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}, \tag{7}$$

yielding

$$[\Sigma_1'] = \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix}, \qquad [\Sigma_2'] = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}. \tag{8}$$

Note that although this is similar to the Burau representation of the Artin braid group for $n = 3$, Eq. (2), it is not a valid representation of that group since $\Sigma_1'\Sigma_2'\Sigma_1' \neq \Sigma_2'\Sigma_1'\Sigma_2'$.

Now if we define

$$H = [\Sigma_2'] = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}, \qquad V = [\Sigma_1'^{\,-1}] = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}, \tag{9}$$

we have recovered two matrices that are identical in form to H and V in D'Alessandro *et al.* [33], with their parameter $a = 2$. They proved that the sequence that maximises the dilatation for a product of matrices from the set $\{H, V\}$, beginning with H, is $HVHVHV\ldots$. (For a sequence beginning with H, multiplication by H^{-1} or V^{-1} never increases the entropy as much as multiplication by H or V, basically because H^{-1} and V^{-1} involve negative matrix elements.) Thus, the optimal protocol with these generators consists of repeating

$$HV = \begin{pmatrix} 5 & 2 \\ 2 & 1 \end{pmatrix}, \tag{10}$$

which has dilatation $3 + 2\sqrt{2} = (1 + \sqrt{2})^2 = \phi_2^2$. As mentioned in Section IV B, the number $\phi_2 = 1+\sqrt{2}$ is the silver ratio. For this reason, we call the braid $\Sigma_2\Sigma_1^{-1}$ corresponding to HV the *silver braid*.

Figure 7 shows a different way of looking at the action of the braid $\Sigma_2\Sigma_1^{-1}$. Here we represent the two rods of Fig. 6 on a periodic lattice, which is topologically equivalent to the annulus. (We can also regard the periodic lattice as the universal cover of the annulus.) In

Fig. 7(b) we see the action of Σ_2 on a material line, followed by the action of Σ_1^{-1} in Fig. 7(c). If we repeat the process, then the length of the material line is stretched asymptotically by at least a factor of ϕ_2^2 per application of $\Sigma_2\Sigma_1^{-1}$. Observe that we can implement a four-rod protocol with the same dilatation by simply moving two pairs of rods at once, a six-rod protocol by moving three pairs, etc. It is this scalability with number of rods that makes the annular braids attractive: the optimal solution can be implemented with any even number of rods.

Another great advantage of this scenario is that the silver ratio $\phi_2 \simeq 2.4142$ is considerably greater than the golden ratio $\phi_1 \simeq 1.6180$, so that material lines are stretched much faster (at almost twice the rate). A final desirable feature is that an annular configuration is natural from an engineering perspective, as we will see in Section VI.

Note that the silver braid having greater entropy per generator than the golden braid does not contradict the optimality conjecture of Section IV, since that applied to a bounded domain, whereas here we have a periodic array of rods. We have also examined topological mixing in periodic and biperiodic geometries in [13].

VI. SILVER MIXERS

The great advantage of the configuration of Section V is that these optimal silver braids are readily implemented with rotating machinery, despite the apparently complicated braiding motion. The easiest way to do this is by using moving planetary gears orbiting a central fixed gear. A horizontal arm and stirring rod attached beneath each planetary gear traces out a cycloid pattern. Within this pattern fixed baffles can be placed, and if the gear ratios and arm lengths are chosen appropriately the moving rods will execute the over-and-under motion with the fixed baffles and produce the silver braid (Section V). We call *silver mixers* stirring devices whose entropy is a multiple of $\log \phi_2$. The taffy-pulling device of Fig. 2 is an example of a silver mixer.

To illustrate the ease of construction, in Fig. 8 we depict gear and arm arrangements which produce silver braids using ordinary gearing available in toy Technic Lego^{TM}. The red (dark) rods are attached to the moving planetary gears, while the blue (light) ones are fixed. A central rod is also employed, indicated in grey. Often the central rod plays no significant topological role, and can be removed. A good reason for retaining the central rod, however, is that it can support a lower set of arms that hold the fixed rods, as in Fig. 9. This convenient feature means that all of the mixing apparatus is in a single unit, and it is not necessary to use a custom-made mixing vessel with baffles fixed to the base (as for the mixograph in Fig. 1). (Kobayashi and Umeda [36] have also constructed a pseudo-Anosov stirring device out of blocks, but it is not optimal.)

We show only the first six in the family of silver mixers, but in principle one could produce the corresponding mixers with $14, 16, \ldots$ rods (plus the central rod). Such devices would be impractical, needing large numbers of very small gears and requiring small arms for the moving rods to prevent them from colliding with each other.

In Fig. 10 we present results from simulating passive advection of a material line in the devices in Fig. 8 in the Stokes flow regime, using the usual particle insertion algorithms to maintain good resolution. (The initial condition in each case being a vertical line from the top of the mixer to the bottom.) Each simulation ran until exponential stretching overwhelmed computational resources, which in these optimised devices happens very quickly. For instance, in Fig. 10(f) we were able to simulate only one complete orbit of the planetary

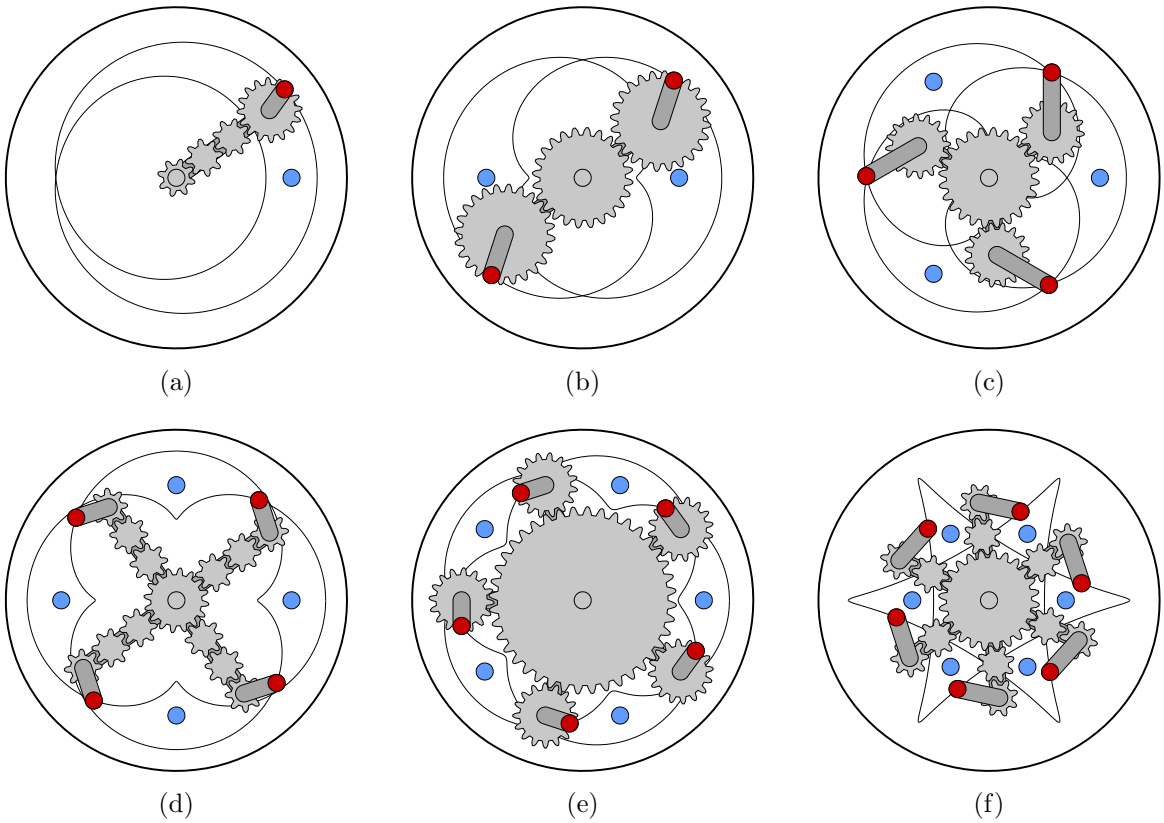

FIG. 8: The first few silver mixers from two (top-left) to twelve (bottom-right) rods, excluding the central rod. The paths of the moving rods are shown, as well as the arms connecting them to the planetary gears. The device in (a) is topologically identical to the taffy-puller of Fig. 2.

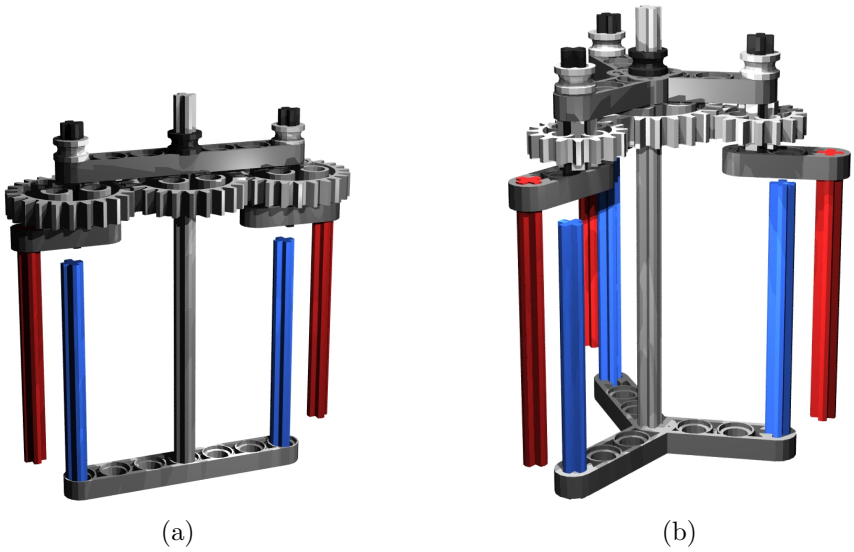

FIG. 9: Implementation of the silver mixing protocols in Figs. 8(b) and 8(c) using Lego^{TM} pieces [35].

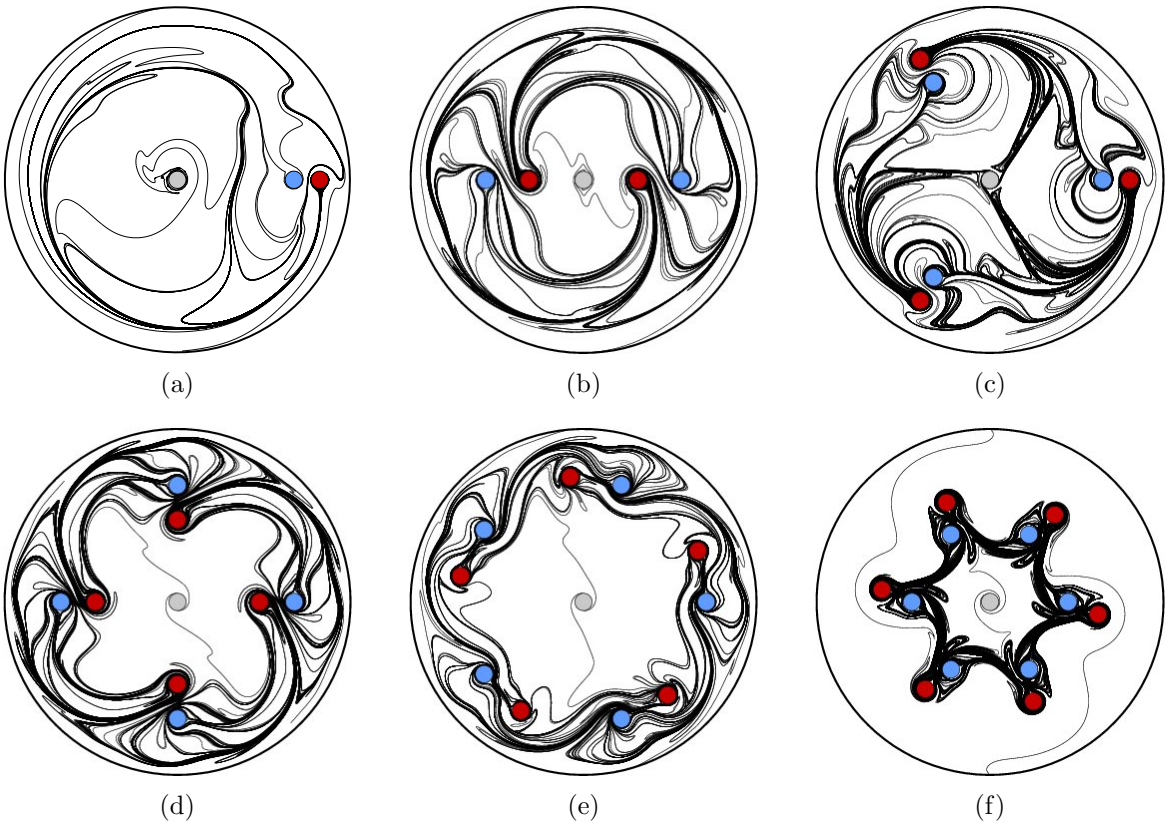

FIG. 10: Mixing patterns for the first few silver mixers, corresponding to the panels in Fig. 8.

gears around the central rod.

If there are n moving rods and n fixed outer rods then one complete orbit of a planetary gear around the fixed gear corresponds to $n/2$ motions $\Sigma_2\Sigma_1^{-1}$, and material lines are stretched by a factor of $(1+\sqrt{2})^{n/2}$. (A silver braid is completed every two orbits when n is odd.) If the period of the flow is defined by a single orbit of the planetary gears then the entropy can be increased by including further pairs of rods, but only at the expense of more complicated machinery and the additional energy input associated with lots of rods moving in close proximity.

Of course, topological entropy is not the only consideration when designing a mixing device. A large mixing region is also desirable, but to determine the size of the mixing region one must solve the particular dynamical equations governing the flow. The exact dimensions of the gears, arms and mixing rods will also have a significant effect on other mixing measures, and these could be tuned by further simulation, provided that the apparatus still produces a silver braid. We have made no attempt to optimise other mixing measures here, since this introduces a wide range of other factors depending on the specific application. We note that in each simulation in Fig. 10 the region of good stretching is commensurate with the extent of the paths of the rods, as is common in Stokes flow mixers [14].

We close this section by presenting a comparison of simulations and experimental results for a silver mixer similar to that in Figs. 8(b) and 9(a). For convenience, in this implementation we arranged for all four outer rods to rotate (see Fig. 11). The two planetary rods rotate in the same direction in small circles around the centre of the planetary gears, while

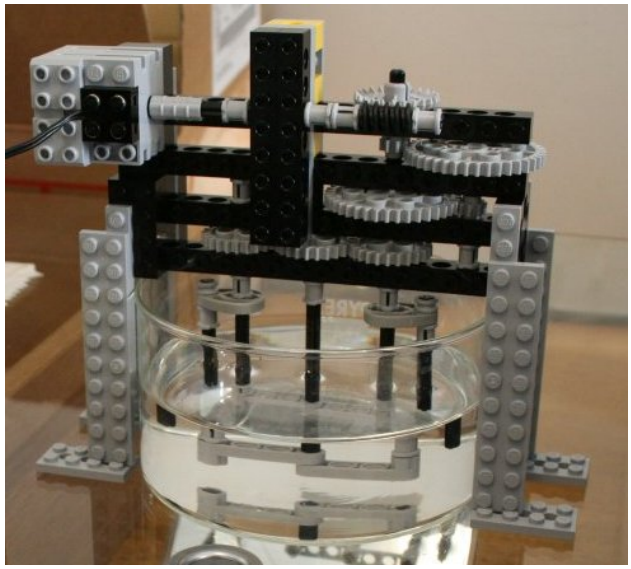

FIG. 11: Photograph of the experimental apparatus, corresponding to Figs. 8(b) and 9(a) in a rotating frame of reference in which all the rods are moving. Note that the extra gearing above the row of central and planetary gears is to reduce the speed of rotation to obtain slow (Stokes) flow.

the two 'fixed' rods rotate around the central rod in the opposite direction, supported from below by arms attached to the central rod. In a rotating frame of reference this motion is identical to that in Fig. 8(b), except for an additional rotation of the outer boundary of the mixing vessel.

The fluid used for the experiments was viscous golden syrup, and the Stokes flow regime was attained by gearing down the motor so that a single rotation of the rods would take several minutes (giving a Reynolds number of order 10^{-4}). We used golden syrup mixed with black food colouring as passive tracer, and tracked the evolution of a circular blob over four complete rotations of the central rod. The blob was photographed from below every half-rotation through the glass worktop using a 45-degree mirror, as shown in Fig. 12.

There is remarkably good agreement between the experimental pictures and the corresponding numerical simulations, with surface and bottom effects causing little discrepancy from the idealised two-dimensional flow. It is readily observed that while the stretching is very rapid, as it must be, the mixing pattern is qualitatively different from that in Fig. 10(b). This is due to the additional rotation of the mixer boundary in a frame of reference rotating with the central rod, which in this case appears to significantly reduce the extent of the region in which the rapid stretching is observed. Nonetheless, the same high stretching rate is found in both mixers, as guaranteed by the braiding motion.

VII. DISCUSSION

In this paper, we discussed the optimisation of a common measure of mixing quality—the topological entropy. This is a crude measure based on the motion of stirring rods, but it is relevant to many fluid-dynamical situations, and crucial for pulling materials such as bread dough or taffy. The entropy is basically the rate of growth of a hypothetical piece of dough

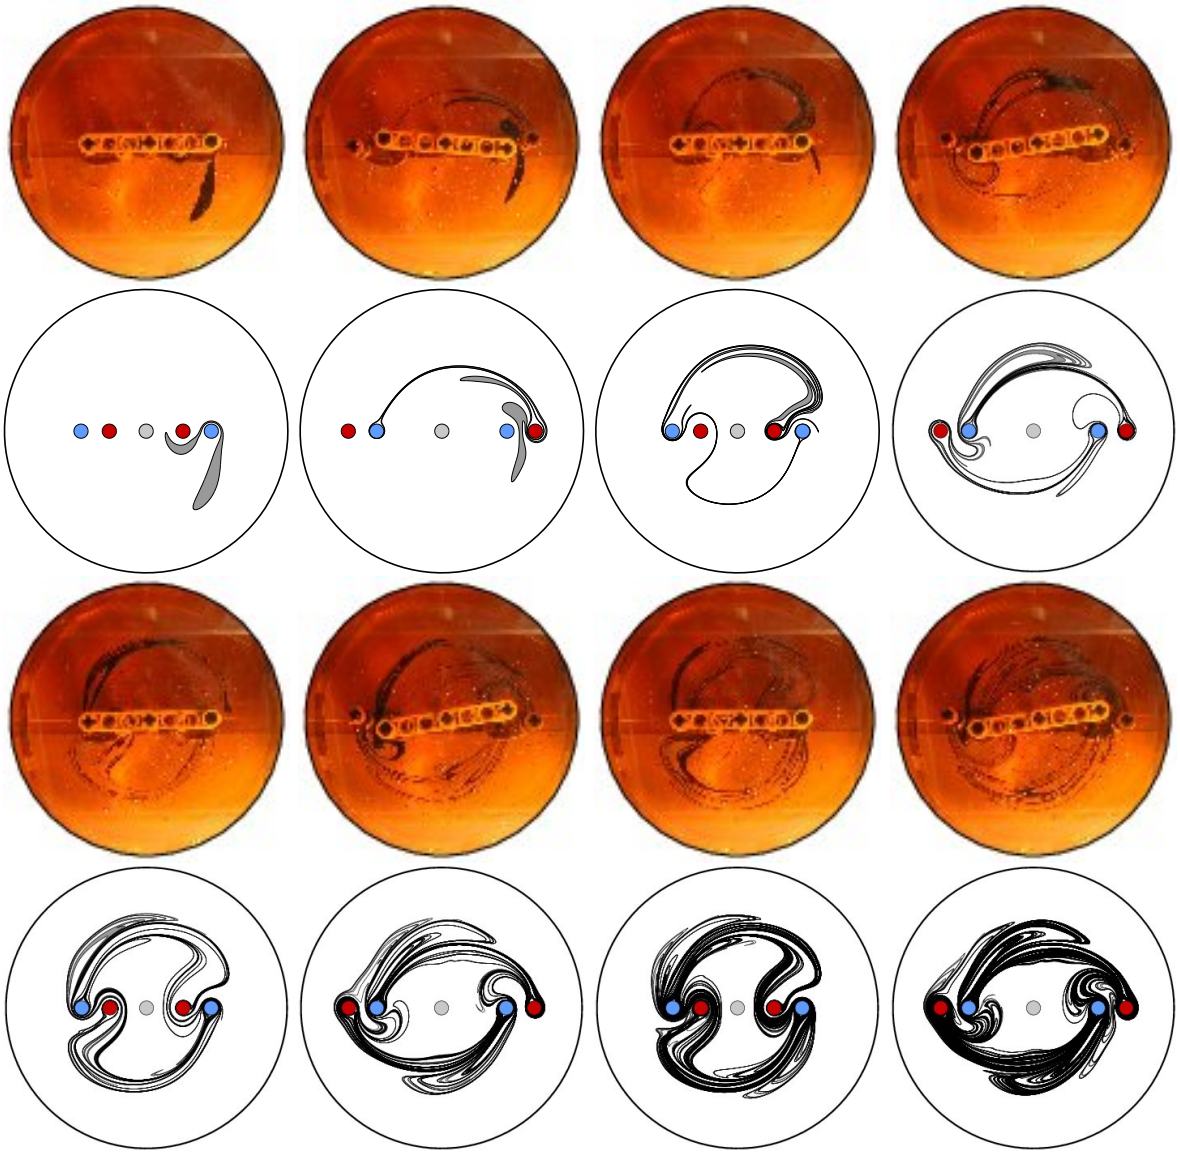

FIG. 12: Comparison of experimental and computational snapshots for the 4-rod silver mixer. The patterns are virtually identical, indicating that three-dimensional effects are negligible, in spite of the supporting horizontal Lego rod spinning at the bottom of the mixing vessel.

wrapped around the rods.

Rod motions are best expressed in terms of the standard generators of the braid group, which allows us to assign a 'cost' to rod motion for optimisation purposes. We examined two such cost-normalised functions: the topological entropy per generator (TEPG) and the topological entropy per operation (TEPO). The latter is less susceptible to assigning high cost to relatively simple motions. For a large number of rods, the TEPG converges to the logarithm of the golden ratio, and the TEPO converge to the logarithm of the silver ratio, though we do not have rigorous proofs of these facts.

We then discussed optimisation using generators of the braid group on the annulus for two rods, which is the natural limit of the TEPO for an infinite number of rods. This spatially-

periodic configuration lends itself well to mechanical construction. We demonstrated this by giving an explicit construction using simple gearing of stirring devices with entropy given by multiples of $\log \phi_2 = \log(1+\sqrt{2})$, where ϕ_2 is the silver ratio. We also built such a device out of Lego^{TM} pieces, as a proof-of-concept.

The implementation of the silver mixers with planetary gearing closely resembles the mixograph discussed earlier in Section I and shown in Fig. 1. However, the mixograph has four moving rods but only three fixed baffles, so the rod motion lacks the rotational symmetry seen in the optimised silver braids. It is instructive to compare the stretching produced by the mixograph with the silver mixer shown in Fig. 9(b) (since this device also has seven rods, including the central one).

To make a fair comparison we note that the braid describing the mixograph in Fig. 4(b) is based on the shortest time required for the rods to return to the same configuration in the rotating frame of reference. This braid has an entropy of approximately 1.4317. Now consider the silver mixer in a rotating frame where the rods on planetary gears rotate in circles and the 'fixed' rods counter-rotate. When the rods return to their initial configuration exactly one silver braid has been completed, with entropy 1.7628. This is about 23% greater than the mixograph. In this sense the mixograph is suboptimal, though the broken symmetry of the rod arrangement improves other aspects of mixing, such as uniformity. For future work it would be desirable to define more 'universal' cost functions that are either less dependent on the specific geometry, or instead better suited to particular engineering applications.

Acknowledgements

We thank Phil Boyland for helping us with the proof sketch in the appendix, as well as Kiran Desai, Emmanuelle Gouillart, Toby Hall, and Jacques-Olivier Moussafir for fruitful discussions. J-LT was partially supported by the Division of Mathematical Sciences of the US National Science Foundation, under grant DMS-0806821. MDF was supported by the Australian Research Council, under Discovery Project DP0881054.

Appendix A: Braids with Metallic Mean Dilatation

From (2), the Burau representation for the braid word $\sigma_1^{-m}\sigma_2^m$ is

$$\sigma_1^{-m}\sigma_2^m = \begin{pmatrix} 1 & m \\ m & 1+m^2 \end{pmatrix} \tag{A1}$$

with dilatation ϕ_m^2, where ϕ_m are the so-called *metallic means* [37]

$$\phi_m := \tfrac{1}{2}(m + \sqrt{m^2+4})\,. \tag{A2}$$

ϕ_1 is the golden ratio, and ϕ_2 is known as the silver ratio. There is a simple geometrical construction of the metallic means: start with a rectangle with one side of unit length, and remove m unit squares. The ratio of the sides of the remaining rectangle is given by the mth metallic mean if it is the same ratio as the original rectangle. The metallic means have continued fraction representation $\{m, m, m, m, \ldots\}$.

The three-rod protocol with dilatation ϕ_m^2 corresponds to m anticlockwise interchanges of rods 1 and 2, followed by m clockwise interchanges of rods 2 and 3. Of course, not all braids

with metallic mean dilatation are of the form $\sigma_1^{-m}\sigma_2^m$; for instance, the annular braid $\Sigma_1^{-1}\Sigma_2$ (Eq. (10)) has entropy equal to ϕ_2^2.

Appendix B: Optimal Topological Entropy: Sketch of Proof

Phil Boyland suggested that the methods used to obtain an upper bound on the efficiency in his unpublished work on π_1-protocols would also work for the topological entropy per generator (TEPG, see Section IV A). (The π_1 stirring protocols have a single moving rod, with all the other rods fixed; they are called π_1-protocols because their rod motions correspond to generators of the fundamental group.) We sketch a proof here, and extend it to the TEPO.

Consider the action of a braid group generator σ_i ($i = 1, \ldots, n-1$) on the standard generators x_j ($j = 0, \ldots, n$) of the fundamental group π_1 of the disc with n punctures. (We include a generator x_0 around the puncture corresponding to the disc's boundary when viewed as a sphere.) This action is given by [21]

$$\begin{aligned} x_i &\mapsto x_i^{-1} x_{i+1} x_i \,, \\ x_{i+1} &\mapsto x_i \,, \\ x_j &\mapsto x_j \,, \qquad\qquad \text{if } j \neq i \text{ or } i+1. \end{aligned}$$

Alternatively, we define the generators $y_j = x_j \cdots x_0$, on which σ_i acts as

$$y_i \mapsto y_{i-1}\, y_i^{-1}\, y_{i+1} \,, \tag{B1a}$$

$$y_j \mapsto y_j \,, \qquad\qquad \text{if } j \neq i. \tag{B1b}$$

The incidence matrix which just counts word length without cancellations for σ_i is called K_i. This matrix has ones on the diagonal and zeros elsewhere (from (B1b)), except $(K_i)_{i,i-1} = (K_i)_{i,i+1} = 1$ (from (B1a)).

For pseudo-Anosov mappings, the topological entropy is equal to the growth rate of word length in π_1 [19, 38] under the action of the mapping. This rate bounded above by the growth of iterates of the corresponding products of incidence matrices. The maximal growth per generator is thus the log of the 'joint spectral radius' (JSR) [39] of the set of matrices $\mathcal{K} = \{K_i : i = 1, \ldots, n-1\}$. We then have, for a given n,

$$\text{TEPG} \leq \log \text{JSR}(\mathcal{K}) \leq \log \max_i \|K_i\|,$$

where $\|\cdot\|$ is any induced norm. Since the 1-norm of all the K_i is 2, it immediately follows that TEPG $\leq \log 2 \simeq 0.6931$. We can improve this by finding the exact JSR of the matrices. We first define the sup norm over all products in $\mathcal{K}$ of length k,

$$\hat{\rho}_k(\mathcal{K}, \|\cdot\|) = \sup \left\{ \left\| \prod_{m=1}^{k} M^{(m)} \right\| : M^{(m)} \in \mathcal{K} \right\}$$

and use the relation

$$\text{JSR}(\mathcal{K}) = \limsup_{k\to\infty} \left(\hat{\rho}_k(\mathcal{K}, \|\cdot\|)\right)^{1/k}$$

which holds for bounded sets of matrices [40]. We will use the 1-norm $\|\cdot\|_1$, which is the sup over column-sums. The form of the matrices K_i means that a matrix $M^{(k)}$ with column-sums $c_j^{(k)}$, $j = 0, \ldots, n$, will after right-multiplication by K_i have new column-sums

$$\begin{aligned} c_{i-1}^{(k+1)} &= c_{i-1}^{(k)} + c_i^{(k)}\,, \\ c_{i+1}^{(k+1)} &= c_i^{(k)} + c_{i+1}^{(k)}\,, \\ c_j^{(k+1)} &= c_j^{(k)}\,, \qquad\qquad j \neq i-1 \text{ or } i+1. \end{aligned}$$

If we then right-multiply by K_{i+1}, we have

$$\begin{aligned} c_i^{(k+2)} &= c_i^{(k+1)} + c_{i+1}^{(k+1)}\,, \\ c_{i+2}^{(k+2)} &= c_{i+1}^{(k+1)} + c_{i+2}^{(k+1)}\,, \\ c_j^{(k+2)} &= c_j^{(k+1)}\,, \qquad\qquad j \neq i \text{ or } i+2. \end{aligned}$$

We can rewrite the net result of K_iK_{i+1} on columns i, $i+1$ as

$$\begin{aligned} c_i^{(k+2)} &= c_{i+1}^{(k+1)} + c_i^{(k)}\,, \\ c_{i+1}^{(k+2)} &= c_i^{(k+1)} + c_{i+1}^{(k)}\,, \end{aligned}$$

where we used $c_i^{(k+1)} = c_i^{(k)}$. Now if we define the sequence $a^{(k)} = c_i^{(k)}$ for k even, and $a^{(k)} = c_{i+1}^{(k)}$ for k odd, then $a^{(k+2)} = a^{(k+1)} + a^{(k)}$, the Fibonacci sequence. If the initial matrix $M^{(0)}$ is the identity, then $a^{(k)} = F_{k+2}$, where F_k is the kth Fibonacci number. This is the largest of the column-sums, so $\hat{\rho}_k(\mathcal{K}, \|\cdot\|_1) = F_{k+2}$. Inserting a product by any other matrix but K_i, K_{i+1} cannot cause the largest column-sum to increase faster, since it will involve the sum of elements that have not been added together. (To complete the proof we also need to show that repeated products such as K_i^2 have a smaller specral radius than K_iK_{i+1}, which is straightforward.) Hence, we have $\hat{\rho}_k(\mathcal{K}, \|\cdot\|_1) = F_{k+2}$, so that $\mathrm{JSR}(\mathcal{K}) = \limsup_{k\to\infty} F_{k+2}^{1/k} = \phi_1$, the golden ratio. Since we have an explicit braid realising this TEPG for $n = 3$ and 4, we conclude that the optimal TEPG is equal to the golden ratio.

For $n > 4$, we can approach golden ratio TEPG as close as we want by using very long braids (see Section IV A); however, we have not yet shown that there are *no* braids realising a golden ratio TEPG for $n > 4$. To show this, observe that the earlier sequence $F_{k+2}^{1/k}$ approaches the golden ratio from above, which is needed for an actual golden ratio TEPG to exist for *finite* k, since $\mathrm{JSR}(\mathcal{K}) \leq (\hat{\rho}_k(\mathcal{K}, \|\cdot\|_1))^{1/k}$ for any k. If we add the constraint of using all the generators, which is a necessary condition for irreducibility (and hence for the mapping to be pseudo-Anosov), then we 'delay' the sequence by at least one term, and it now approaches the golden ratio from below (see Fig. 13). Hence, for finite k the JSR can never be equal to the golden ratio. This means that for $n > 4$ we have TEPG $< \phi_1$, strictly.

To find an upper bound on the topological entropy per operation (TEPO, see Section IV B), we inflate the set $\{K_i\}$ into the set $\mathcal{L} = \{L_j\}$ of all the operations that can be performed simultaneously (i.e., products of K_i that all commute, without repeating any K_i). For example, for $n = 6$ the set $\mathcal{L}$ consists of 12 matrices:

$$\mathcal{L} = \mathcal{K} \cup \{K_1K_3, K_1K_4, K_1K_5, K_2K_4, K_2K_5, K_3K_5\} \cup \{K_1K_3K_5\}$$

In general, the set $\mathcal{L}$ has cardinality $\#\{L_j\} = \sum_{j=1}^{n} \binom{n-j}{j} = F_{n+1} - 1$, where F_n it the nth Fibonacci number. Since the largest column-sum in $\mathcal{L}$ is always 3, we immediately get

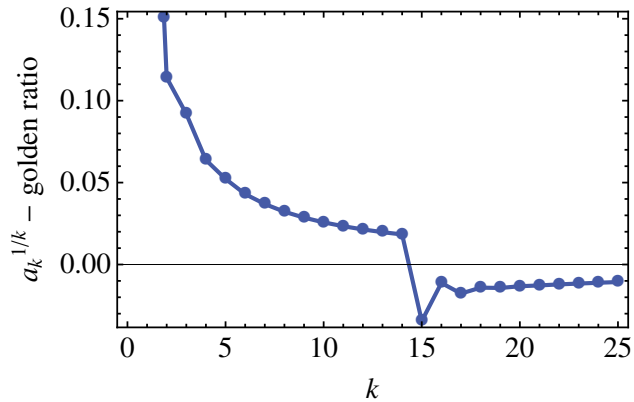

FIG. 13: Convergence of the sequence $a_k^{1/k}$ to the golden ratio: here $a_{k+2} = a_{k+1} + a_k$, except for $a_{15} = a_{14}$. The sequence approaches the golden ratio from below after the skipped term.

the bound TEPO $\leq \log 3$ independent of n. If we use the 2-norm instead, we get the result TEPO $< \log(1+\sqrt{2})$ (the inequality in this case is strict), again independent of n. Since the TEPO of the sequence (4) asymptotes to $\log(1+\sqrt{2})$ for large n, then the optimal TEPO converges to $\log(1+\sqrt{2})$ as claimed in Section IV B.

Intuition suggests that the largest joint spectral radius in $\mathcal{L}$ is given by the product $L_{\text{odd}}L_{\text{even}}$, with

$$L_{\text{odd}} = K_1K_3K_5\cdots, \qquad L_{\text{even}} = K_2K_4K_6\cdots.$$

In fact the product $L_{\text{odd}}L_{\text{even}}$ gives exactly the same TEPO as the sequence (4) for each n. Hence, proving that the JSR of the set L is realised by the product $L_{\text{odd}}L_{\text{even}}$ for all n is enough to prove the optimality of (4). Proving that we have the JSR appears much harder than for the TEPG.

Note that it seems fortuitous that for both the optimal TEPG and TEPO the upper bound given by the joint spectral radius of the incidence matrices is sharp. If this were not the case, then proving that we have the TEPG or TEPO would involve more that just linear algebra techniques. Intuitively, one can argue that the braids realising the optimal TEPG and TEPO avoid 'cancellations' in growth of words in π_1, which is not surprising since such cancellations would make the braid less efficient.

---

[1] P. L. Boyland, H. Aref, and M. A. Stremler, 'Topological fluid mechanics of stirring,' *J. Fluid Mech.* **403**, 277–304 (2000).

[2] B. J. Binder and S. M. Cox, 'A mixer design for the pigtail braid,' *Fluid Dyn. Res.* **40**, 34–44 (2008).

[3] P. L. Boyland, M. A. Stremler, and H. Aref, 'Topological fluid mechanics of point vortex motions,' *Physica D* **175**, 69–95 (2003).

[4] E. Gouillart, M. D. Finn, and J.-L. Thiffeault, 'Topological mixing with ghost rods,' *Phys. Rev. E* **73**, 036311 (2006).

[5] J.-L. Thiffeault and M. D. Finn, 'Topology, braids, and mixing in fluids,' *Phil. Trans. R. Soc. Lond. A* **364**, 3251–3266 (2006).

[6] J.-L. Thiffeault, M. D. Finn, E. Gouillart, and T. Hall, 'Topology of chaotic mixing patterns,' *Chaos* **18**, 033123 (2008).

[7] J.-L. Thiffeault, 'Braids of entangled particle trajectories,' *Chaos* **20**, 017516 (2010).

[8] P. L. Boyland, 'Topological methods in surface dynamics,' *Topology Appl.* **58**, 223–298 (1994).

[9] J.-L. Thiffeault, 'Measuring topological chaos,' *Phys. Rev. Lett.* **94** (8), 084502 (2005).

[10] J.-L. Thiffeault, E. Gouillart, and M. D. Finn, 'The size of ghost rods,' in L. Cortelezzi and I. Mezić, eds., *Analysis and Control of Mixing with Applications to Micro and Macro Flow Processes*, *CISM International Centre for Mechanical Sciences*, volume 510, 339–350 (Springer, Vienna, 2009), `http://arxiv.org/abs/nlin/0510076`.

[11] A. Vikhansky, 'Simulation of topological chaos in laminar flows,' *Chaos* **14** (1), 14–22 (2004).

[12] M. A. Stremler and J. Chen, 'Generating topological chaos in lid-driven cavity flow,' *Phys. Fluids* **19**, 103602 (2007).

[13] M. D. Finn, J.-L. Thiffeault, and E. Gouillart, 'Topological chaos in spatially periodic mixers,' *Physica D* **221** (1), 92–100 (2006).

[14] M. D. Finn, S. M. Cox, and H. M. Byrne, 'Topological chaos in inviscid and viscous mixers,' *J. Fluid Mech.* **493**, 345–361 (2003).

[15] M. D. Finn, S. M. Cox, and H. M. Byrne, 'Chaotic advection in a braided pipe mixer,' *Phys. Fluids* **15** (11), L77–L80 (2003).

[16] J.-O. Moussafir, 'On computing the entropy of braids,' *Func. Anal. and Other Math.* **1** (1), 37–46 (2006).

[17] M. D. Finn and J.-L. Thiffeault, 'Topological entropy of braids on the torus,' *SIAM J. Appl. Dyn. Sys.* **6**, 79–98 (2007).

[18] R. K. Connelly and J. Valenti-Jordan, 'Mixing analysis of a Newtonian fluid in a 3D planetary pin mixer,' *Chem. Eng. Res. Design* **86** (12), 1434–1440 (2008).

[19] A. Fathi, F. Laundenbach, and V. Poénaru, 'Travaux de Thurston sur les surfaces,' *Astérisque* **66-67**, 1–284 (1979).

[20] W. P. Thurston, 'On the geometry and dynamics of diffeomorphisms of surfaces,' *Bull. Am. Math. Soc.* **19**, 417–431 (1988).

[21] J. S. Birman, *Braids, Links, and Mapping Class Groups*, Annals of Mathematics Studies (Princeton University Press, Princeton, NJ, 1975).

[22] E. Artin, 'Theory of braids,' *Ann. Math.* **48** (1), 101–126 (1947).

[23] S. E. Newhouse and T. Pignataro, 'On the estimation of topological entropy,' *J. Stat. Phys.* **72** (5-6), 1331–1351 (1993).

[24] Y. Yomdin, 'Volume growth and entropy,' *Israel J. Math.* **57** (3), 285–300 (1987).

[25] S. E. Newhouse, 'Entropy and volume,' *Ergod. Th. Dynam. Sys.* **8**, 283–299 (1988).

[26] J.-L. Thiffeault, E. Lanneau, and S. E. Matz, 'The cat's cradle, stirring, and topological complexity,' Dynamical Systems Magazine, `http://www.dynamicalsystems.org/ma/ma/display?item=292`.

[27] W. Burau, 'Über Zopfgruppen und gleichsinnig verdrilte Verkettungen,' *Abh. Math. Semin. Hamburg Univ.* **11**, 171–178 (1936).

[28] B. Kolev, 'Entropie topologique et représentation de Burau,' *C. R. Acad. Sci. Sér. I* **309** (13), 835–838 (1989), english translation at `http://arxiv.org/abs/math.DS/0304105`.

[29] G. Band and P. L. Boyland, 'The Burau estimate for the entropy of a braid,' *Algeb. Geom. Topology* **7**, 1345–1378 (2007).

[30] D. Fried, 'Entropy and twisted cohomology,' *Topology* **25** (4), 455–470 (1986).

[31] M. Bestvina and M. Handel, 'Train-tracks for surface homeomorphisms,' *Topology* **34** (1),

109–140 (1995).

[32] T. Hall, '*Train: A C++ program for computing train tracks of surface homeomorphisms*,' `http://www.liv.ac.uk/maths/PURE/MIN_SET/CONTENT/members/T_Hall.html`.

[33] D. D'Alessandro, M. Dahleh, and I. Mezić, 'Control of mixing in fluid flow: A maximum entropy approach,' *IEEE Transactions on Automatic Control* **44** (10), 1852–1863 (1999).

[34] V. D. Blondel, Y. Nesterov, and J. Theys, 'On the accuracy of the ellipsoid norm approximation of the joint spectral radius,' *Linear Algebra Appl.* **394**, 91–107 (2005).

[35] These images were drawn with LeoCAD, a freely available Lego CAD program (`http://leocad.org/`), and rendered with PoVRay, the freely available Persistence of Vision Raytracer software (`http://www.povray.org/`).

[36] T. Kobayashi and S. Umeda, 'Realizing pseudo-Anosov egg beaters with simple mecanisms,' in *Proceedings of the International Workshop on Knot Theory for Scientific Objects, Osaka, Japan*, 97–109 (Osaka Municipal Universities Press, 2007).

[37] E. W. Weisstein, 'Silver ratio,' From *MathWorld*—A Wolfram Web Resource. `http://mathworld.wolfram.com/SilverRatio.html`.

[38] R. Bowen, 'Entropy and the fundamental group,' in *Structure of Attractors*, *Lecture Notes in Math.*, volume 668, 21–29 (Springer, New York, 1978).

[39] G.-C. Rota and G. W. Strang, 'A note on the joint spectral radius,' *Nederl. Akad. Wet., Proc., Ser. A* **63**, 379–381 (1960).

[40] M. A. Berger and Y. Wang, 'Bounded semigroups of matrices,' *Linear Algebra Appl.* **166**, 21–27 (1992).